SIMULATING INITIAL CONDITIONS IN AGENT-BASED MODELING

W.M. BULLEIT* and M.W. DREWEK, Michigan Tech, Houghton, MI

ABSTRACT

Predicting where terrorists are most likely to strike concerns planners, law enforcement and government agencies at various levels, and engineers who must design facilities of all kinds. The present work is an effort to use agent-based modeling to examine the interaction of civilians, terrorists, and security to determine the types of facilities in a town or city that are most susceptible to attack. Agent modeling of civil violence has been performed in the past. The ultimate goal of our research is to be able to estimate the probability of attack for various types of facilities in a population center so that resources can be allocated for hardening or otherwise protecting those facilities. Because of the nature of resource-based agent modeling, the agents must be allowed to evolve in the town or city environment before the day-to-day behavior of the community is simulated. We have approached that problem by breaking the total simulation into two parts: (1) the incubation of the community, where the agent population evolves to live in the environment, and (2) the simulation of the behavior of the evolved agents in the community environment. Results from this work indicate that incubation can be ended at any desired time and still allow modified time-step simulation. This result allows modified time-step simulation of a population in any stage of its evolution. When transitioning from incubation to simulation, the behavior of the population must be allowed to stabilize in the early stages of the shortened time-step simulation.

Keywords: Agent-based modeling, artificial societies, simulation, terrorism

INTRODUCTION

Predicting where terrorists are most likely to strike concerns planners, law enforcement and government agencies at various levels, and engineers who must design facilities of all kinds. The present work is an effort to use agent-based modeling to examine the interaction of civilians, terrorists, and security to determine the types of facilities that are most susceptible to attack. Agent modeling of civil violence has been performed in the past (Epstein 2002). The ultimate goal of our research is to be able to estimate the probability of attack for various types of facilities in a population center so that resources can be allocated for hardening or otherwise protecting those facilities.

Agent models comprise a range of types, of which this one is an extension of the type used by Epstein and Axtell (1996) in which society evolves by using the basic concepts of resources in the environment, agent metabolism for those resources, and agent vision (knowledge of the environment). This model represents a community in which civilians evolve to become radicals (inactive terrorists) who may become active terrorists committing attacks on the community. The environment in which this community evolves consists of a rectangular grid on

---

* *Corresponding author address*: William M. Bulleit, Department of Civil and Environmental Engineering, Michigan Tech, 1400 Townsend Dr., Houghton, MI 49931; email: wmbullei@mtu.edu.

which a number of resources lie. Civilian agents evolve in this environment on the basis of their vision and metabolism for the various resources. The terrorist agents evolve from the civilian agents by using a tag-mediated procedure derived from that used by Axelrod (1997). After the agent becomes a terrorist, it remains an inactive terrorist until its age and wealth each reach a specific value that allows it to become active. An active terrorist agent stops looking for resources and begins to examine the agent wealth within its vision. When it finds a location of high local wealth, it moves to that location and becomes a suicide bomber that explodes, destroying the agents and the wealth on the surrounding grid points. The number and location of security agents are determined on the basis of the wealth, fear, and innate nervousness of the agents in the civilian population. The number of security agents in the community evolves as attacks occur. Security agents search for and arrest terrorists in regions of locally high populations.

Results from this basic model (Bulleit and Drewek 2005) show that the location of attacks is affected by the choice of the base level of security. Higher base levels of security shift many of the attacks away from the areas of high resources. In this work, a base level of security does not exist; security levels are endogenous. Thus, it appears that endogenous agent modeling of communities will require the use of an *incubation* period during which the community can evolve to allow the agents to acclimate to the environment and develop a set of initial conditions that are themselves endogenous. A limited use of incubation has been used by Cederman (2003). In that case, he merely allowed the simulation to run for a set number of time-steps before beginning data collection. We propose a more distinct incubation period. In the proposed incubation period, the time-step will be longer than what will later be used for the community simulation from which results will be gleaned. For instance, during the incubation period, the time-step might be representative of a year. The community will be allowed to evolve during the incubation to a user-chosen time. At that point, the time-step will be shortened (e.g., to a day), and the simulation will continue with the conditions at the end of incubation becoming the initial conditions for the short time-step simulation.

The objective of this paper is to describe the use of a simulation process that has an *incubator* in which the community evolves to a certain point and a *simulator* in which the day-to-day community simulation is performed.

MODEL DESCRIPTION

Community Environment

The environment in which this community evolves consists of a 50×50 rectangular grid on which lie a number of *piles* of resources. Each civilian agent requires a set amount of each different resource. The resource piles can be isolated in the sense that there may not be a resource gradient between the piles. This lack of gradient is important to the design of the civilian agents. For this study, the environment consists of four resource piles, each representing a different resource. All agents require each resource to live. A second aspect of the environment relates to the effects of a terrorist attack on the environment. A terrorist attack, modeled as a suicide bomber, results in the destruction of all resources on the grid point where the terrorist was at the time of the attack plus all agents, all their wealth, and all resources on the Moore neighborhood of that grid point. The resources at these nodes remain zero for 2 years before they begin to regenerate. The attack also makes agents *fear* the grid points where the attack occurred. The level of fear that agents feel for the attacked nodes dissipates with time and spreads to surrounding nodes. Figure 1 shows von Neumann and Moore neighborhoods.

For the environment that we discuss in this paper, the range of resource values for each resource is a maximum of 54.0 units and a minimum of 1.0 unit. The grow-back rates in the incubator are one-fourth of the maximum value allowed at each node. Hence, the maximum grow back rate is 13.5 units/year, and the minimum is 0.25 unit/year. Figure 2 shows the environment. The maximum resources are on the *peaks,* and the minimums are on the *plains*.

Civilian Agents

Civilian agents evolve in the environment. Each agent is assigned an initial metabolism for each different resource in the environment from a uniform distribution with a range of 1.5 to 3.0 or U(1.5, 3.0) for resources 1 and 3 and U(1.25, 2.50) for resources 2 and 4. The initial agent vision is an integer selected from U(3, 7). Vision is the number of grid points that an agent can see in the four cardinal directions from its current location. The agents are also randomly assigned an amount of each of the different resources in the environment from U(45, 90), their *wealth.* Thus, an agent's wealth is an agent's store of each of the various resources in resource units. Each agent's *initial endowment* is randomly selected from U(12, 24), in units of *generalized resources*. *A* generalized resource for an agent is one of its resources divided by the metabolism for that resource, thereby converting resource units into a time or, in other words, the amount of time an agent can live, assuming that it collects no more of the given resource. Initial endowment is discussed subsequently. The agents' death age is an integer selected randomly from U(40, 80), and the agent's nervousness factor is randomly selected from U(0, 1). Nervousness is a measure of how nervous an agent is in the presence of fear. Last, each agent in

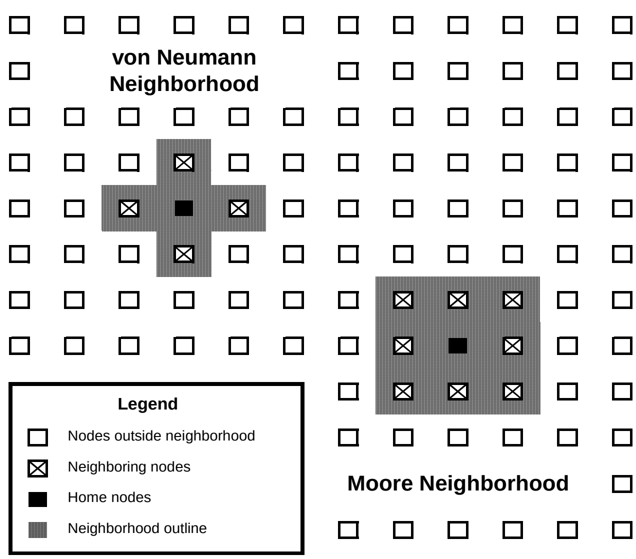

FIGURE 1 Von Neumann and Moore neighborhoods

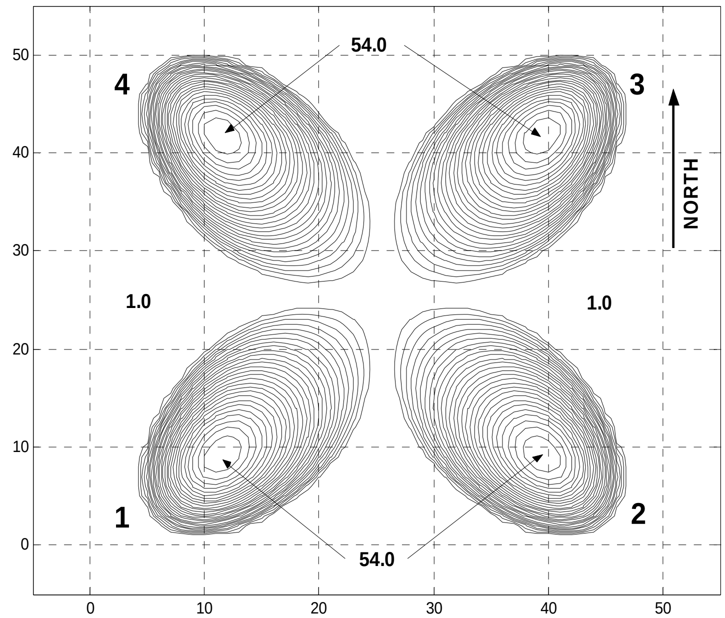

FIGURE 2 Community environment

the initial population is given a cultural tag in which each of the five tag integers is randomly selected from U(0, 9).

The initial population is made up of agents of an age between 10 and 50. Since there is some overlap in the initial age range and death age range, if the death age selected is less than the initial age, another death age is selected until the death age is greater than or equal to the initial age. Once the agents' initial wealth has been determined, the agents' initial *generalized wealth* is calculated. Generalized wealth is the length of time an agent can live assuming it collects no additional resources of any kind; i.e., the minimum of the generalized resources.

The agents move around the environment in search of resources that they need to live. The agents search for their *critical resource*. The critical resource is the resource currently limiting the agent's life span, assuming no additional resources are collected. The critical resource is the resource yielding the minimum generalized resource and is the resource used in determining an agent's generalized wealth. Because the environment does not have a resource gradient at all locations, the agents were given memory. Without this memory, it is difficult to evolve a stable population. The resource memory is simple: the agent remembers the grid point where the maximum of each of the different resources that it has encountered in its travels around the environment is located. Thus, in this case, since there are four different resources in the environment, the agent stores the location and amount of the maximum value of each of the four resources it has encountered. It updates these values as it finds a better source (larger value) of a specific resource. As well as allowing a stable population to evolve, this simple memory allows agents to evolve patterns of travel between resource locations; for instance, the path between two resources could represent travel between home and work in a real community.

Agents also have a memory of the maximum and minimum fear they have seen as they traveled around the environment. For the baseline case, discussed subsequently, fear memory includes the past 5 years. As an agent searches its local environment for resources, it considers the fear associated with the nodes it is examining. The maximum and minimum fears seen during a given time-step are recorded and will be remembered for the next 5 years. After 5 years have passed, maximum and minimum fears for the sixth year are forgotten. Fear and its use are described in detail below.

Agents have a gender, and when male and female agents meet they procreate if each of them has reached a fertile age and is wealthy. Procreation allows the agents' vision, metabolism, and nervousness to evolve. A potential parent is an agent that is fertile (i.e., has an age within the fertility age range) and possesses a generalized wealth equal to or greater than its own initial endowment. The minimum fertility age for males is an integer selected randomly from U(12, 15) with the maximum from U(50, 60). For females, the selection is made from U(12, 15) and U(40, 50), respectively. When an agent moves to a node, if that agent is a potential parent and one of its von Neumann neighbors is also a potential parent, and assuming that in one of their von Neumann neighborhoods there is an unoccupied node, then an child is born. If more than one potential parent of opposite gender is located in the agent's von Neumann neighborhood and an unoccupied node is still available, then the mate is selected at random. Potential parent agents who have a parent/child relationship or share a common parent are not allowed to procreate.

When a newborn agent is added to the population, its placement in the environment is selected randomly from all the unoccupied nodes in the parents' von Neumann neighborhoods. The newborn's vision is determined by taking the average of the vision of the parents (rounded to the nearest integer) with a mutation probability, P_{mv}, of 0.0025 that this value will be increased by 1.0 or decreased by 1.0. (Vision is limited to a minimum of zero and has no set maximum.) Infant agent metabolisms are determined in the same way, with mutation probability $P_{mm} = 0.0025$, except that the minimum metabolism cannot drop below the minimum of the range of the uniform distribution used in the selection of metabolisms for the initial agents (i.e., 1.25 or 1.5). Infant agent nervousness is determined the same way, with mutation probability $P_{mn} = 0.0025$, but the change is either +0.1 or –0.1. Nervousness is kept within the range of 0 to 1. The newborn's initial wealth is calculated by multiplying one-half of the father's initial endowment by his metabolism for each resource and adding to that the corresponding results of a similar calculation for the mother. The mother and father each donate the resources to their newborn; the resources donated are forfeited from the parents. This store of individual resources is used to determine the newborn agent's initial generalized wealth by dividing each resource level by the newborn's metabolism for each respective resource. The newborn agent's initial generalized wealth serves as its initial endowment. The newborn's gender is selected at random, with an equal chance of each. The newborn's fertility age range is selected from the ranges used by the initial population, as is the newborn's death age. The newborn's initial knowledge of where resources lie in the environment is taken from each parent's memory: The parents give the newborn the "best" locations of each resource in either of their memories. Note that the parents also exchange the best resource locations in their respective memories. The newborn's cultural tag is determined from its parents' tags; for each tag integer, there is an equal probability that the value will be taken from the mother or the father.

Terrorist Agents and Terrorist Attacks

Terrorist agents evolve from the civilian agent population. The evolution of a civilian agent to a terrorist is performed by using a tag-mediated process that is based on the approach used by Axelrod (1997). As described above, each agent is assigned a tag at the beginning of the simulation or at birth. The tag consists of a string of five integers in which each integer ranges from 0 to 9. As the agents move around, they interact with other agents. The interaction is controlled by the tags, and the evolution of a civilian to a terrorist is based on the tag values. First, consider interaction. When an agent moves to a grid point, it examines, at random, one of the grid points in its von Neumann neighborhood. If an agent is in that location, the agents compare the sum of the absolute value of the difference between each of the five integers in their tag:

$$S = \sum_{i=1}^{5} \left| I_{ij} - I_{ik} \right| , \tag{1}$$

where I is the value of the tag integer, i is the location of the tag integer, and j and k are the indexes of the interacting agents. The larger this sum is, the smaller the probability that the agents interact. If the sum is 45, then the probability is 0.0 that they interact. If the sum is 0, then the probability of interaction is 1.0. The probability of interaction is linear between these two end points. If the agents interact, then one of the integer locations on the tag is chosen at random — a 0.20 probability that any one of the five is chosen. Once one of the integer locations is chosen, the agents compare the integer they have at that location. If the integers are the same, nothing happens. If the integers are different, then one of two things occurs: (1) the agent that moved changes its integer to match the agent that it interacted with, or (2) the agent that moved has a radical change. The probability of a radical change is determined from using:

$$P_{rc} = P_b \frac{\left| I_{ij} - I_{ik} \right|}{9} , \tag{2}$$

where P_{rc} is the probability of radical change, P_b is the base probability (a P_b of 0.02 is used in all example simulations), and I_{ij} and I_{ik} have been defined previously. The direction of the radical change is determined by using the *changing* agent's current integer value. For example, if the current integer is 2, then there is a 2/9 probability that the agent will change to a 9, and a 7/9 probability that the value will change to 0. The agent that moved will be the changing agent. After the agents have interacted, whether or not an integer change has occurred in either of the above two ways, there is still a small, isolated change probability, P_{ic}, of 0.02 that one of the integers on its tag will change by –1 or +1. This ends the interaction. The agent that moves has the changes occur to it so that there is no possibility that an agent will be changed more than once during any time-step (Axelrod 1997).

An agent becomes a terrorist on the basis of the sum of the five integers in its tag (referred to as *cultural identity*). The probability that the agent becomes a terrorist is determined by using a U-shaped symmetrical polynomial function that passes through 1.0 at a sum of 0, passes through 0.0 at a sum of 22.5, and passes through 1.0 again at a sum of 45. Figure 3 shows the U-shaped curve. Thus, there is some probability that any agent can become a terrorist, but the probability is greatest near the end points of the sum of the tag integers. After the agent becomes a terrorist, it remains an *inactive* terrorist until its age and wealth each reach a specific value that allows it to become active. An inactive terrorist agent becomes *active* if it is 18 years old or older and its generalized wealth is equal to 5.0 or greater. Once active, the terrorist agent will remain active as long as its generalized wealth remains greater than 3.0. After every change to the tag,

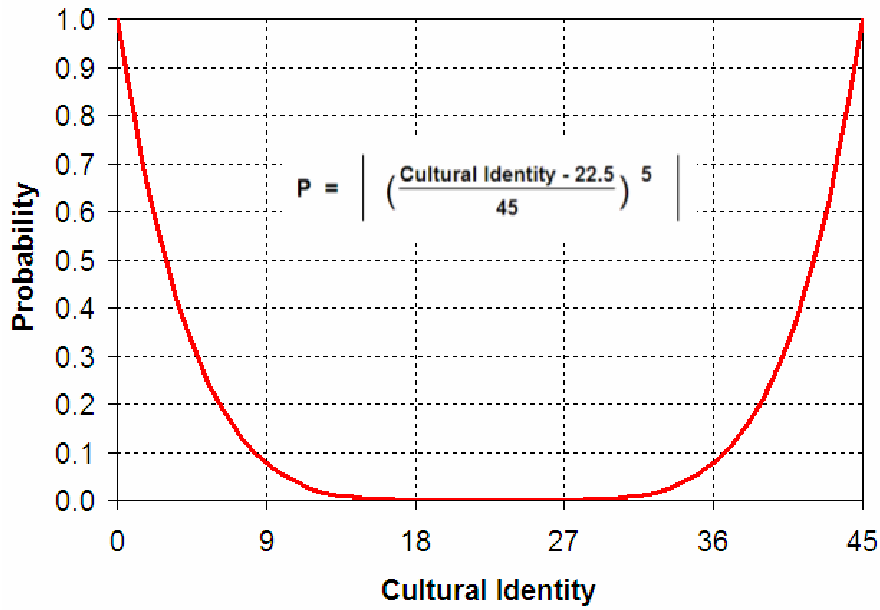

FIGURE 3 Probability of becoming a terrorist

the agent's new sum is used to determine the probability of becoming a terrorist (if the agent is already a terrorist, the probability is that of remaining a terrorist).

An active terrorist agent stops looking for resources and begins to examine the wealth on the von Neumann neighborhood of grid points within its vision and moves to the node with the largest surrounding wealth. (Note that even though the terrorist stops looking for resources, it continues to collect resources at the nodes it is visiting) This wealth information, referred to as surveillance data, consists of the present agent generalized wealth and the moving average of the agent generalized wealth on each grid point in the von Neumann neighborhood over the past 10 time-steps (referred to as *historical nodal wealth*).

This approach is used because terrorists do not strike just very wealthy locations but also locations where wealth passes through (e.g., airports). The active terrorist agent then keeps track of the mean and standard deviation of the largest five surveillance data values that it has seen in its travels. When it finds a grid point that has a surveillance data value that is greater than the mean plus some number of standard deviations (typically 1.0) *and* the coefficient of variation of its surveillance data is less than 0.25, it becomes a suicide bomber and explodes, destroying wealth on the Moore neighborhood as discussed above. These two criteria for detonating allow an active terrorist agent to attack when it finds a local region with a relatively consistent high level of wealth.

When a terrorist agent conducts an attack, all of the agents with their wealth and the nodal resources are destroyed in the terrorist agent's Moore neighborhood. Agents will fear the nodes in the destroyed area and, over time, in areas surrounding the destroyed area. The sum of all the agent generalized wealth destroyed becomes the fear at each of the nodes in that Moore neighborhood. If the level of fear is greater than the level that existed before the attack, then the portion of the fear that is greater than the existing fear level will diffuse outward over time, reducing the nodal intensity of the fear. Eventually, if enough time elapses without another terrorist attack, the fear level from the attack becomes uniform across the environment, thereby affecting every agent equally and thus having no effect on any agent's decision process. The spreading of fear has been modeled in the same way that Epstein and Axtell (1996) modeled the diffusion of pollution> Details on the diffusion of fear can be found in that reference. If the level of fear is less than the level that existed before the attack, then the fear level remains the same as it was before the attack for those nodes, that is, new fear is not summed on top of existing fear.

Agents evaluate the critical resource that they see by using the *adjusted critical resource*. The adjusted critical resource is used to evaluate the resources on a given node. The adjusted critical resource is the amount of the critical resource at a node adjusted to take into account the fear level at that node, the fear memory of the agent, the generalized wealth of the agent, and the innate nervousness of the agent. The resource values (for the critical resource) at each node are adjusted according to the following equation:

$$AR_i = R_i - \Delta R_i \ , \tag{3}$$

where

$$\Delta R_i = R_i n \left[1 - \exp\left(-\frac{(f_i - f_{min})}{(f_{max} - f_{min})} \times \frac{(f_i - f_{min})}{9GW} \right) \right], \tag{4}$$

where AR_i is the adjusted critical resources located at node i, R_i is the amount of the critical resource located at node i, ΔR_i is the amount of resource located at node i that an agent is willing to sacrifice for less fear, n is the agent's innate nervousness factor, f_i is the fear at node i, f_{max} is the largest fear seen in the last 5 years, f_{min} is the smallest fear seen in the last 5 years, and GW is the agent's generalized wealth. The two terms making up the exponent serve two purposes: the first term normalizes the fear level at node i to the range of fear seen in the recent past, and the second term normalizes the relative fear, $f_i - f_{min}$, to the agent's generalized wealth. The "9" appears in the second term because there are nine nodes in a Moore neighborhood. Once each node has been considered, the agent moves to the node with the largest adjusted critical resource.

Security Agents

The number of security agents in the community evolves as attacks occur. The number of security agents is based on characteristics of each agent in the population. These characteristics include the wealth of the agent, the resources that the agent collects at each time-step, the level of fear that the agent feels at that time-step, the maximum and minimum amount of fear that the agent has felt in the past, and the inherent nervousness of the agent. These characteristics are used to determine the amount of resources that the agent is willing to contribute to buying security. Note that the agents do not actually give up any resources. One method of putting a dollar value on a nonmarket good is to conduct a survey, essentially asking people how much they would be willing to pay for something to happen (Dorfman and Dorfman 1993). The responses are summed for the population affected, and this becomes an estimate for the value of that nonmarket good. This is called *contingent evaluation* and corresponds to the approach we are using to assign security to the environment.

At each time-step, adult civilian agents (agents that have reached their minimum fertility age) consider what portion of the resources being collected at that time-step they would be willing to contribute to purchase security. This willingness to contribute resources, without actually giving them up, is the agent's demand for security. Equation 5 is used to calculate the amount of each resource that each agent is willing to contribute:

$$C_{\mathrm{j}} = \mathrm{R}_{\mathrm{ij}}\mathrm{n}\left[1 - \exp\left(-\frac{(\mathrm{f}_0 - \mathrm{f}_{\mathrm{min}})}{(\mathrm{f}_{\mathrm{max}} - \mathrm{f}_{\mathrm{min}})} \times \frac{(\mathrm{f}_0 - \mathrm{f}_{\mathrm{min}})}{9\mathrm{GW}}\right)\right], \tag{5}$$

where C_j is the contribution of resource j, R_{ij} is the amount of resource j at node i, f_0 is the maximum fear that the agent has seen during that time-step, and the other variables are the same as in Equation 4. The first term making up the exponent in Equation 5 normalizes f_0 to the range of fear seen by the agent in the recent past, and the second term normalizes $f_0 - f_{min}$ to the agent's generalized wealth. During the time-step, the agent contributions are summed, and at the end of the time-step, there is a pool of each resource. The average metabolism for all nonsecurity agents for each of these resources is determined. Each resource pool is divided by the average metabolism for that resource. The minimum of these values becomes the number of security agents required at the end of the time-step (i.e., the number of security agents is the average number that can be supported by the contributed resources).

If the existing number of security agents needs to be increased to meet the calculated requirement, new security agents are introduced in the environment. The new security agents are given vision randomly selected from the range given by the absolute minimum and absolute maximum vision in the agent population. Nodes with higher historical nodal wealth have a higher probability of receiving these new security agents. Specifically, each unoccupied node is assigned a random number from U(0,1). Each of these random numbers is multiplied by the historical nodal wealth at the node divided by the maximum historical nodal wealth found in the environment. After all of the unoccupied nodes have been considered, the new agents are located on the nodes with the largest adjusted random number.

If the existing number of security agents needs to be decreased to meet the requirement, some existing security agents are removed from the environment. The security agents located on nodes with lower historical nodal wealth have a higher probability of being removed. The process for removing security agents is the same as for adding them, except that the security agents located on the nodes with the smallest adjusted random number are removed from the environment.

Security agents search for terrorists in regions of high population. Each security agent moves to the open grid point within its vision that has the most agents on its von Neumann neighborhood. Once on that grid point, the security agent examines its von Neumann neighborhood. It interacts with (investigates) each agent on the von Neumann neighborhood with a probability related to the number of agents in the neighborhood; for example, if there are three agents in the von Neumann neighborhood, then it interacts with each of those agents with a probability of one-third. If the security agent interacts with an agent, there are two possible outcomes: (1) it releases civilians or (2) it arrests terrorists (active or inactive), with a probability determined by using the U-shaped symmetrical polynomial function described above. When used for security agents arresting terrorist agents, the U-shaped symmetrical polynomial function is cubic. (When used for generating terrorist agents, the U-shaped symmetrical polynomial function is fifth order, as shown in Figure 3.) The probability of arresting a terrorist agent increases for more radical agents (sum of tag integers closer to 0 or 45). An arrested terrorist agent is permanently removed from the environment.

Incubation and Community Simulation

The day-to-day behavior of the community becomes apparent only after the agents, particularly the civilian agents, have learned to live in the environment. The *incubation period* is that time during which the agents are allowed to adjust to the environment. The portion of the simulation encompassing the incubation period is referred to as the *incubator.* The time-step of the incubator is representative of a *year*. The portion of the simulation following the incubator is referred to as the *simulator.* When switching from the incubator to the simulator, the model must be calibrated to the adjusted time-step. First, the number of time-steps that make up a year, T_y, must be chosen. At each time-step, the current time of the simulation is incremented by $1/T_y$. For ease of discussion, we refer to a time-step with a duration of $1/T_y$ as a *day.* (If the incubation time-step is representative of a year, then using $T_y = 365$ would produce the day that we are familiar with.) All time-related parameters must be adjusted. The agents' age, in years, is converted from an integer to a real number by adding a random number from U(0,1). In the simulator, each agent's age is incremented at each time-step by $1/T_y$. The agents' maximum and minimum fertility ages, as well as the agents' death ages, remain integers. To maintain consistent agent evaluations of resources, in order to maintain stable agent wealth when switching between the incubator and simulator, the resource concentrations in the environment are divided by T_y. The agents' resource metabolisms are also divided by T_y.

When a terrorist attack occurs, all resources are destroyed on the nodes involved, and the area remains devoid of all resources for 2 years. In the incubator, this time period is equal to two time-steps. In the simulator, the damaged area also remains devoid of all resources for 2 years, but the number of time-steps is $2 * T_y$. Fear is generated in the same way in the incubator and simulator, but the fear dispersion rate (on a per-time-step basis) must be adjusted. In the incubator, fear dispersion on a per-time-step basis is the rate at which fear disperses in a year. In the simulator, on the first day, the fear dispersion that will occur over the first year is calculated by using the procedure from the incubator. The resulting change in fear over the next year at each node is then divided by T_y, producing the change in fear per day at each node. After 1 year passes or when a terrorist attack occurs or when 1 year has passed since the last terrorist attack, the change in fear per day at each node is recalculated.

The occurrence probability for procreation and cultural exchange are reduced from 1.0 in the incubator to $1/T_y$ in the simulator. This modification means that in the incubator, a potential parent agent will consider procreating at each time-step, but in the simulator, this will happen at each time-step with a reduced probability. The same is true for cultural exchanges. In the incubator, an agent will consider culturally interacting with a neighbor at each time-step. In the simulator, at each time-step, an agent will consider culturally interacting with only one of its neighbors, with a probability of $1/T_y$.

Agent resource memory remains unchanged, except that the resource magnitudes held in memory are divided by T_y. Agent fear memory also changes in the transition from the incubator to the simulator. For the baseline case, fear memory consists of 5 years of memory of the largest and smallest nodal fears seen each year. In the incubator, this means keeping track of the largest and smallest nodal fears seen at each time-step. However, in the simulator, 1 year is composed of multiple time-steps. At the end of 1 year, an agent remembers the minimum and maximum fear seen in the past year, t_1, as well as the minimum and maximum fear seen in the previous four years, t_2, t_3, t_4, and t_5. At the beginning of the new year, minimum and maximum fears seen in t_5 are forgotten. The present year becomes t_1; what had been referred to as t_1 becomes t_2, t_2 becomes t_3, and so on. For the first day of the new year t_1, the minimum and maximum fear remembered are the minimum and maximum fear seen while searching the environment on that day. On the second day of year t_1, the minimum fear seen is compared to the minimum fear remembered on the first day, and the smallest value is remembered. A similar comparison is done for maximum fear. The procedure is repeated for each day in t_1. When the year is over, at the beginning of a new year, the minimum and maximum fear remembered in t_5 is forgotten, and the process is repeated.

The last change in the transition from the incubator to the simulator involves the determination of agent historical nodal wealth. For the baseline case, as well as all sensitivity studies done for this paper, the historical nodal wealth is the average agent generalized wealth that has been present on a node over the previous 10 years. In the incubator, this is easily calculated, since 10 years equal 10 time-steps. However, in the simulator, determining historical nodal wealth is not so easy. When an agent first moves into the simulator, it has 10 years' worth of data from the incubator and nothing from the simulator. On the first day of a year in the simulator, the agent generalized wealth present on each node that day is determined, and the value at each node is divided by T_y. Since the historical nodal wealth is calculated by using 10 years' worth of data, on the first day of a year in the simulator, the data would be taken from the first day of year t_1; all of the data collected for years t_2 through t_{10}; and 364 days out of 365 for year t_{11}. In this way, the historical nodal wealth is still the average nodal wealth over a 10-year period. In general, for any day in year t_1, the historical nodal wealth at a given node can be calculated from the equation:

$$HNW_{1,k} = \frac{1}{10}\left[\sum_{x=1}^{k} GW_{1,x} + GW_2 + \ldots + GW_{10} + \left(\frac{T_y - k}{T_k}\right) GW_{11}\right], \qquad (6)$$

where $HNW_{1,k}$ is the historical nodal wealth calculated on day k of year t_1, $GW_{1,x}$ is the agent generalized wealth present on the node on day x of year t_1, GW_2 is the total agent generalized wealth present on the node over year t_2 (similar for GW_3 through GW_{11}), and T_y is the number of days in a year. On the last day in the first year of the simulator, year t_1's contribution to the historical nodal wealth is based on T_y days, or one full year, and year t_{11}'s contribution has shrunk to zero. On the first day of the next year in the simulator, all of the year subscripts are incremented by adding one, and year t_1 once again represents the current year, and year t_{12} is forgotten.

RESULTS AND DISCUSSION

The process described above was implemented by using MatLab (MathWorks 2002).

Generating Initial Conditions with Various Incubation Cut-off Times

The simulation of initial conditions involves a two-step process. First, pre-incubation conditions are formulated by using input parameters, where some input parameters define deterministic characteristics of either the environment or the agent population and others define ranges for uniform distributions. The pre-incubation conditions are then used to begin incubation where the time-step is analogous to 1 year. Upon termination of this incubation at a specific time, the ending conditions of both the environment and agent population are recorded. The post-incubation conditions eliminate much of the bias introduced by the user input parameters and the methods used to generate the pre-incubation conditions. More significantly, the post-incubation conditions will also typically represent an agent population acclimatized to its environment. The agents have had time to evolve and gain knowledge of their environment. The post-incubation conditions are then used to begin a simulation in which the time-step is much shorter. The occurrences during these simulations are the occurrences of interest.

In performing this process and analyzing the results, two scenarios were considered. First, a single set of pre-incubation conditions was generated. The input parameters are those defined throughout the previous sections of this paper. These conditions were then used to begin an incubation run. From this incubation run, post-incubation conditions were recorded after 200, 700, and 1,200 time-steps (years). Figure 4 shows the total nonsecurity agent population over the time this incubation run was performed.

In each case, when the post-incubation conditions were generated after 200, 700, and 1,200 years of incubation, they were used as the initial conditions for the simulator. Each simulator run used 365 time-steps per year. Figure 5 shows the total non-security agent population over a 5-year period for the cases where the initial conditions are based on 200-, 700-, and 1,200-year incubations, respectively.

Qualitatively, the simulated population histories continue on from the point at which the incubator left off. For example, when the post-incubation conditions were based on an incubation run of 200 years (see Figure 4), the incubator showed a relatively small population (less than 100 non-security agents) with a relatively small growth rate. In the corresponding simulator population history (see Figure 5), the population continues to be low and the growth rate continues to be small. The opposite is the case in which the post-incubation conditions were based on an incubation run of 700 years. Here the incubator had a moderate population (over 300 nonsecurity agents) and was experiencing rapid growth. In the corresponding simulator population history (see Figure 5), the population is moderate and the growth rate continues to be rather rapid. In the case in which the post-incubation conditions were based on an incubation run of 1,200 years (see Figure 5), the population is large (over 600 non-security agents) and rather stable but becomes somewhat cyclic. At the time that post-incubation conditions were generated, the population was climbing toward the upper cusp of one of those cycles. Not unexpectedly, the simulator shows a large population with a moderate growth rate.

At 200-, 700-, and 1,200-year incubation times, the growth rates shown in Figure 5 are superimposed on Figure 4. In each case, the growth rate from the simulator was significantly greater than the growth trend seen in the incubator in the same time frame. Although the causes of this phenomenon require further experimentation, several observations can be made at this time. When 1,200-year incubation was used, the corresponding simulator growth rate closely

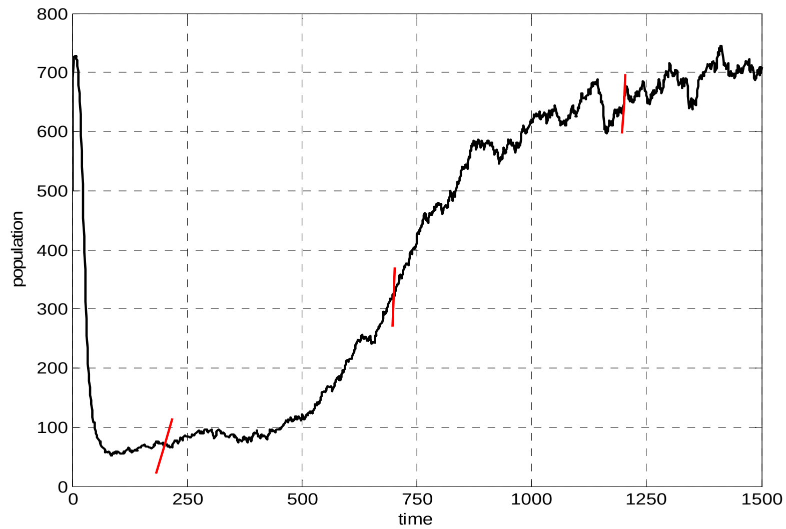

FIGURE 4 Incubator population history over 1,500 years

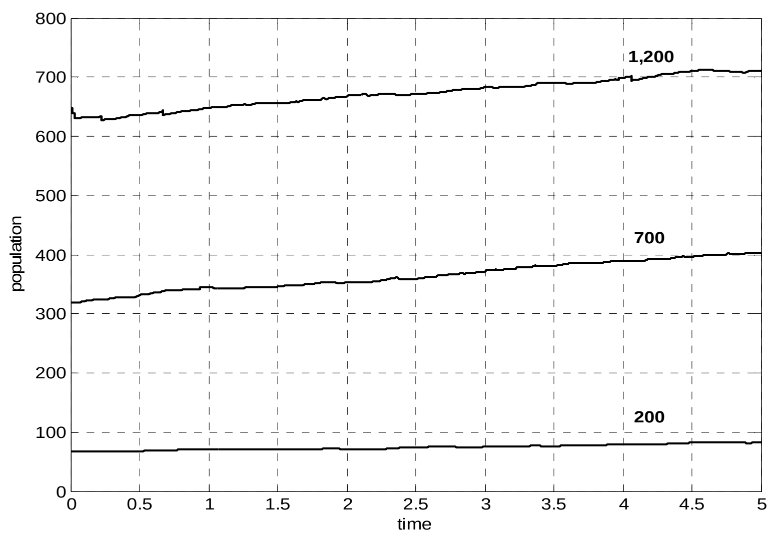

FIGURE 5 Simulator population history from using initial conditions from 200-, 700-, and 1,200-year incubations

matched the growth rate seen in the incubator at approximately 1,200 years over the 5-year period considered. Additional experiments using 5-year simulations in the simulator show that as T_y approaches 1.0, the growth rates in the simulator approach the growth trends in the incubator, as would be expected. Other experiments using 20-year simulations in the simulator with a T_y of 12 and 52 show that the growth rates in the simulator approach the respective growth trends exhibited in the incubator.

The behavior of the total nonsecurity agent population when shifting from incubator to simulator is important to consider. Major changes are neither expected nor desired. What happens in the incubator occurs at a particular rate per year; in the simulator, that particular rate per year is expected to fall within the same range as it does in the incubator, except one year is divided into multiple time-steps. In essence, the incubator speeds through history; the simulator moves slowly from the present into the future. With regard to population (and total population is indicative of many aspects of individual agent behavior), this transition is smooth and appears to be insensitive to the time in the incubation when the transition takes place.

Of course, the population history is not the only comparison of interest. A large amount of data are extracted from both the incubation and simulator runs. For example, the mean cultural identity can be plotted for both the incubation and simulator. The mean cultural identity coming out of the incubator remains relatively stable throughout the simulator. The same is the case for the variation in the cultural identity. The portions of the population with a cultural identity between 0–9, 10–18, 19–27, 28–36, and 37–45 were also examined. The population coming out of the incubator had cultural identity demographics very similar to those throughout the simulator. Other agent population characteristics where also examined, including: wealth, age, vision, resource metabolisms, innate nervousness, and initial endowment. The agent population characteristics at the end of the incubator were similar to the characteristics throughout the 5-year simulator run regardless of the incubation time. For example, when incubation was terminated after 1,200 time-steps, the average agent vision was 10 nodes in the four cardinal directions. The maximum vision in the population was 12; the minimum vision was 9. Throughout the simulation, the maximum and minimum remained the same, although the average vision showed a very slight increase. This behavior was typical of the other characteristics defining the population, including: wealth, resource metabolism, age, death age, and innate nervousness. For the 200-year incubation, more changes were observed in the simulator. The population was smaller; consequently, births and deaths had a larger impact. Similar behavior was observed for the 700-year incubation. The fairly limited adjustments between the incubator and the simulator do not drastically affect the agent population.

The next issue to be examined is the effect on agent behavior as the agents move around the environment. To examine this, consider the historical nodal wealth averaged over the last 10 years of the incubator — specifically the case in which the incubator was run for 1,200 years. Figure 6 is the historical nodal wealth contour plot, showing the average agent generalized wealth present on a node over the last 10 years of the incubation. After the 5-year simulation, another historical nodal wealth contour plot was generated. In this case, the historical nodal wealth comprises the last 5 years of the incubator plus the additional 5 years of the simulator (i.e., it is still based on 10 years of data). Figure 7 shows historical nodal wealth after simulation.

A comparison of Figures 6 and 7 shows no notable differences between the two. Some minor changes in magnitude and variations in contour shapes may exist, but, for the most part, a detailed comparison indicates agent behavior is unchanged between the incubator and simulator. This is especially true for areas where the historical nodal wealth is relatively high and less true for areas where the historical nodal wealth is relatively low. In the simulator, agents are maintaining their wealth and moving around the environment in much the same way as they did in the incubator. A similar comparison was done for population density contours, where the average nodal population over identical 10-year periods was considered. The results of that analysis showed that nodal population density also showed few significant changes between the incubator and simulator.

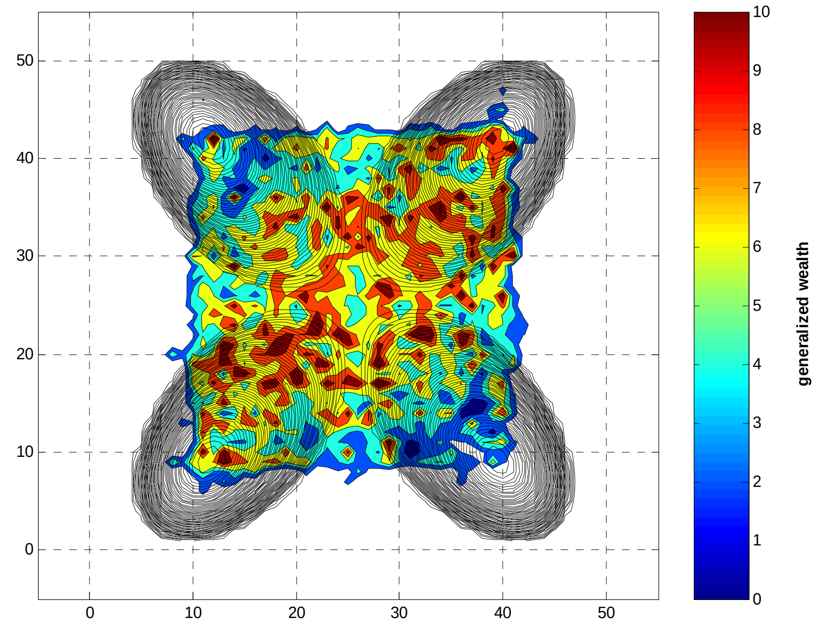

FIGURE 6 Historical nodal wealth over the last 10 years of a 1,200-year incubation

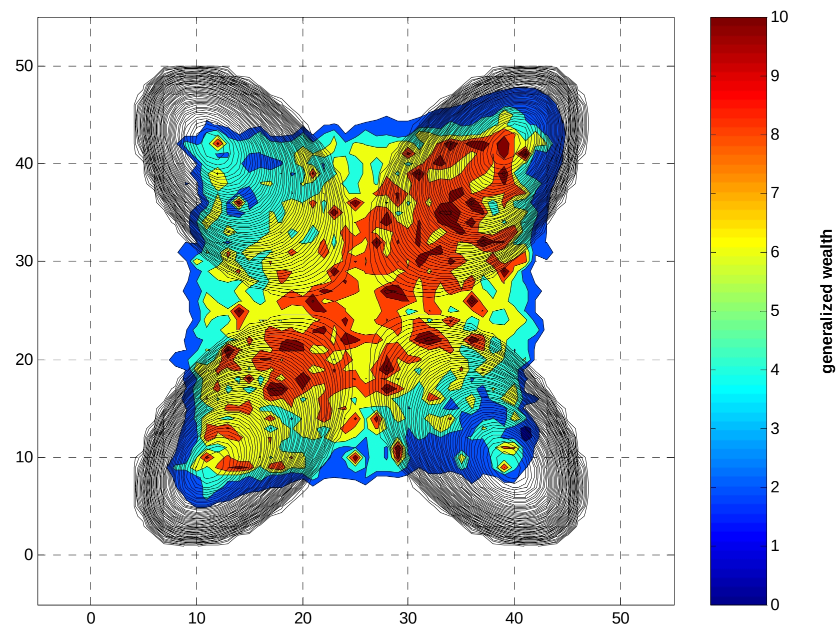

FIGURE 7 Historical nodal wealth after 5 years in the simulator

The formation and behavior of terrorist agents, as well as the generation of security agents within the population, varied significantly among the three scenarios considered here. When the incubation was stopped after 200 years and the simulation began, the population was small, and relatively little terrorist activity had occurred in that short period of instability. Low levels of terrorist activity led to only a small or nonexistent security population. After the transition from incubator to simulator, no additional terrorist activity occurred, and only minimal security was present over short periods of time (a result of fear created by the attacks in the incubator). When the incubator was run for 700 and 1,200 years, the populations were either growing rapidly or were relatively stable, respectively. Terrorist activity within the incubator resulted in a constant security presence by the end of the incubation runs. After the transition from incubation to simulation, the trends appeared to continue. In both cases, multiple terrorist attacks occurred in the 5-year simulations, and security levels jumped upward just after the attacks and slowly declined in the time following. When only the transition between incubation and simulation is considered, the behavior related to terrorism and security appears to be relatively constant even after the time-step definitions change.

However, the transition between incubation and simulation does not appear to be seamless. While the agent demographics, terrorist and security behavior, and aggregate behavior of the simulator remain fairly steady, there is a short time in the beginning of the simulator run in which the behaviors of agents seeking resources change. When an agent seeks its critical resource, it has a tendency to try to *equalize* its generalized resources. If one generalized resource drops below the others, the agent begins searching for the corresponding resource. Provided that the agent remembers where to find adequate concentrations of all the resources to ensure survival, the agent's generalized resources will reach an equilibrium in which each generalized resource is "close" to being equal. How close together the generalized resources can get depends on the magnitudes of the resources collected during a time-step. Since in order to maintain wealth, environmental resource concentrations have been factored by $1/T_y$ in the simulator, when the agent population moves into the simulator, there is a period of time during which an agent will attempt to equalize its generalized resources. Since an agent can collect only a small amount of a given resource at a time-step in the simulator (whereas in the incubator, it can collect 1 year's worth of that resource), the agent has a tendency to spend more time seeking its critical resource. In other words, the agent will spend more time on the pile of its critical resource. Since all agents are doing the same for their critical resources, a noticeable change in behavior can be observed in the early phase of a simulator. This behavior will continue to occur until the differences between each agent's generalized resources are reduced. Once the generalized resources have been equalized, the agent resource usage will once again stabilize. The agent's behavior when seeking resources will again look like it did in the incubator.

The time required for the simulator to stabilize varies with how stable the agent population is at the time of simulation. For example, when the incubator ended after 200 years, the simulator required approximately 1 year to stabilize; in this case, the population is small but there is little growth. When the incubator ended after 700 years, the simulator required approximately 2 years to stabilize; in this case, the population is established but experiencing rapid growth. And when the incubator ended after 1,200 years, the simulator required as little as one-half year to stabilize; in this case, the population is large and relatively stable. Thus, the change in the behavior of agents seeking resources always occurs, no matter how long the incubator is allowed to run. However, the more stable the population, the quicker the simulator stabilizes.

Incubation Runs Using Identical Input Parameters

The next issue to examine is the variability of post-incubation conditions generated by using one set of input parameters and their effects on 5-year simulator runs. For this analysis, four post-incubation conditions were generated, again by using the input parameters discussed previously, and the post-incubation conditions were based on 1,200 years of incubation. The only difference in each of the trial runs was produced by using a different random number seed to generate the initial population. After the post-incubation conditions were generated, each was used to begin a single 5-year simulation, where one year was subdivided into 365 time-steps.

Each of the four incubation runs was successful. A stable population evolved in the environment, and after 1,200 years, the post-incubation nonsecurity agent populations were all about 600 to 700. (Incubation run 4 had an ending population just below 600. This lower population was the result of an increased level of terrorist activity.) The population history for the first incubation run is shown in Figure 4. The behavior over time was similar for all four runs. While some differences did exist (e.g., in the second incubation run), the population dropped to approximately 25 agents around year 200, and the basic shape of the population history remained the same. In each case, the initial population of 500 agents grew rapidly for a brief period of time, then crashed, struggled to gain a foothold in the environment as the agents gained knowledge of their environment and evolved, and this was followed by rapid growth and finally a decreased growth rate as the population stabilized.

Agent characteristics (including, vision, resource metabolisms, wealth, and innate nervousness) evolved in a similar fashion for all four incubation runs. For example, by the end of 1,200 years, in each case, the average agent vision evolved to include approximately 10 nodes in the four cardinal directions. The average agent vision going into the incubator was 5 nodes in the four cardinal directions. The agent resource metabolisms at the end of the incubator were approximately the same. Agent age demographics were also steady between the four incubation runs. Agent generalized wealth (including the absolute maximum and minimum, as well as the average) was also similar after 1,200 years of incubation. In each case, agent innate nervousness (originally uniformly distributed between 0 and 1) averaged somewhere between 0.4 and 0.6, with the absolute maximum and minimum at +0.1 or –0.1 of the average. Thus, the agent characteristics defining the post-incubation agent population were very similar among the four incubation runs.

Even though most agent characteristics were similar among the four incubation runs, the agent cultural identities differed significantly, and this led to differences in the level of terrorist activity and the corresponding level of security present in the environment. Essentially the incubation runs fell into three categories: (1) runs 1 and 2 had a moderate level of terrorist activity over the 1,200-year incubation (for run 1, 20 terrorist attacks occurred, 11 terrorist agents were arrested, and the security level was approximately 3.6% of the nonsecurity population; for run 2, 28 terrorist attacks occurred, 23 terrorist agents were arrested, and the security level was approximately 5.0% of the nonsecurity population); (2) run 3 had significantly more terrorist activity over the 1,200-year incubation (70 terrorist attacks occurred, 121 terrorist agents were arrested, and the security level was approximately 10.0% of the nonsecurity population); and (3) run 4 showed significantly less terrorist activity over the 1,200-year incubation (11 terrorist attacks occurred, 1 terrorist agent was arrested, and the security level was approximately 4.0% of the nonsecurity population at the end of incubation).

Other information about the four incubation runs casts light on the reason for the differences in the conditions among the different incubation runs. First, the level of terrorist activity is directly related to the agents' cultural identity. For example, in incubation run 4, where the level of terrorist activity was relatively low, at one point during the incubation run (specifically, when the population reached its low point), the agent population was entirely made up of agents with a cultural identity between 19 and 27. While this anomaly was quickly eliminated by the population, by the end of the 1,200-year incubation, the group with a cultural identity between 19 and 27 was still dominant, and the fringe groups (those with cultural identities between 0 and 9 and between 37 and 45) were an extremely small portion of the total population. Likewise, in incubation run 3, the group with a cultural identity of 19–27 made up approximately 40% of the population, which led to significant increases in the groups with cultural identities between 10 and 18 and between 28 and 36 (approximately 20% and 25% of the population, respectively). This difference led to larger-than-usual fringe group populations.

Thus, the differences in the level of terrorist activity are directly attributed to the distribution of cultural identities in the population throughout the incubation run. The distribution of cultural identities is controlled not so much by the randomness of the pre-incubation population but by the randomness during the early "collapse" of that pre-incubation population (Figure 4). When the population crashes, the distribution of pre-incubation cultural identities can be drastically modified. Sometimes the distribution is flattened out, thereby adding to the fringe groups. At other times, the distribution is tightened up, and the population tends toward the median. Once the population begins to grow again, cultural interactions lead to further changes in the agents' cultural identities. It is reasonable to assume that given enough time, the distribution of cultural identity for the individual incubation runs would stabilize. However, since each of the four incubations was run for 1,200 years, different cultural identity distributions resulted.

Generally, when the level of terrorist activity was high at the end of an incubation run, the corresponding level of security was also high, and vice versa. But the correlation between these is not particularly strong. When the case with the least terrorist activity, run 4 (where 11 terrorist attacks and 1 arrest occurred) is considered, the security level of 4.0% coming out of the incubator was approximately the same as that of run 1: 3.6% (where 20 terrorist attacks and 11 arrests occurred). Keep in mind that the level of security is related to the level of fear seen in the environment, as well as the maximum and minimum fear seen over the past 5 years. Thus, sporadic terrorist attacks, which allow time for fear to dissipate, can lead to relatively low security levels. Such is the case for incubation runs 1, 2, and 4. In the case where the level of terrorist activity was substantially higher — incubation run 3 (where 70 terrorist attacks and 121 arrests occurred) — the attacks are no longer sporadic. The fear levels are high and consistently increasing; attacks are occurring throughout the environment and thereby effecting larger portions of the population.

One last observation is about the differences between these four different incubation runs and the resulting post-incubation conditions. Notice that only in incubation run 4 did the security get the better of the terrorists (121 terrorist agents arrested for 70 terrorist attacks). Only when the level of terrorist activity was high were the results of security really felt. Intuitively, this makes sense. When terrorist activity is low and attacks are sporadic, it is difficult to keep a sense of urgency in the population; consequently, the level of security is highly variable (increasing immediately after an attack, decreasing in the times when no attacks occur). This leads to a great disadvantage for the security agents. They are not present in substantial enough numbers to keep the threat level under control; therefore, the terrorist agents have the advantage. If the population is complacent about terrorism, the terrorists will gain an advantage, allowing them to conduct attacks with a much lower risk of being thwarted. However, as the level of terrorist activity increases, the level of security also increases, and the terrorist agents are subjected to a significantly greater probability of being arrested.

Post-incubation Initial Conditions in the Simulator

Each of the four individual sets of post-incubation conditions generated previously for 1,200 years of incubation will be used as the initial conditions for a single 5-year simulation where each year is subdivided into 365 time-steps. As was observed before, the nonsecurity agent population defined by the post-incubation conditions was stable throughout the 5-year simulation, with relatively slow growth. The agents' characteristics also remained stable, including vision, metabolism, age demographics (average, maximum, minimum, and average death age), wealth (average, maximum, and minimum), innate nervousness, and initial endowment (average, maximum, and minimum). The cultural demographics of the population were also stable over the 5-year simulation. The cultural demographics measured were the average cultural identity, variation in cultural identity, and the portions of the nonsecurity population made up of agents with a cultural identity of 0–9, 10–18, 19–27, 28–36, and 37–45.

Considering the level of terrorist activity during 5 years in the simulator, the simulation associated with incubation run 1 had 5 terrorist attacks and 4 terrorist arrests. The level of security ranged between 1.5% and 6.5%, with a ballpark average more than 3% over the 5-year simulation. For the simulation associated with incubation run 2 (see Figure 8a-d), 7 terrorist attacks occurred, and 12 terrorist agents were arrested. The level of security ranged between 1.5% and 6.5% (Figure 8c), with an average of about 3% over the 5-year simulation. For the simulation associated with incubation run 3, 12 terrorist attacks occurred, and 13 terrorist agents were arrested. The level of security started the simulation around 15% and steadily declined to approximately 5%. Attacks at the very end of the incubation and early in the 5-year simulation caused a significant increase in security at the incubation/simulation interface. For the simulation associated with incubation run 4, 8 terrorist attacks occurred, and 5 terrorist agents were arrested. The level of security at the beginning of the simulator was about 4%, fell quickly down to 1%, steadily increased to an average of about 7%, and then declined, ending the simulation with an average of 6.5%.

Comparing the level of terrorist activity in the simulator to the results from the incubation reveals some trends. Incubation runs 1 and 2 exhibited moderate levels of terrorist activity, incubation run 3 showed a high level of terrorist activity, and incubation run 4 showed a minimal level of terrorist activity. Note that when these post-incubation conditions were used to begin a 5-year simulation, the level of terrorist activity roughly corresponded to that in the incubator. The level of terrorist activity experienced in simulation runs 1 and 2 was moderate (with 7 and 5 attacks, respectively; averaging 6 attacks in 5 years). Simulation run 3 experienced twice the number of attacks (12 attacks in 5 years). Simulation 4 had 8 attacks in 5 years. At the end of incubation run 4, the number of terrorist attacks was increasing, and the level of security was relatively low. Therefore, the post-incubation conditions depicted a population in a very different phase of its development than in the other three simulations.

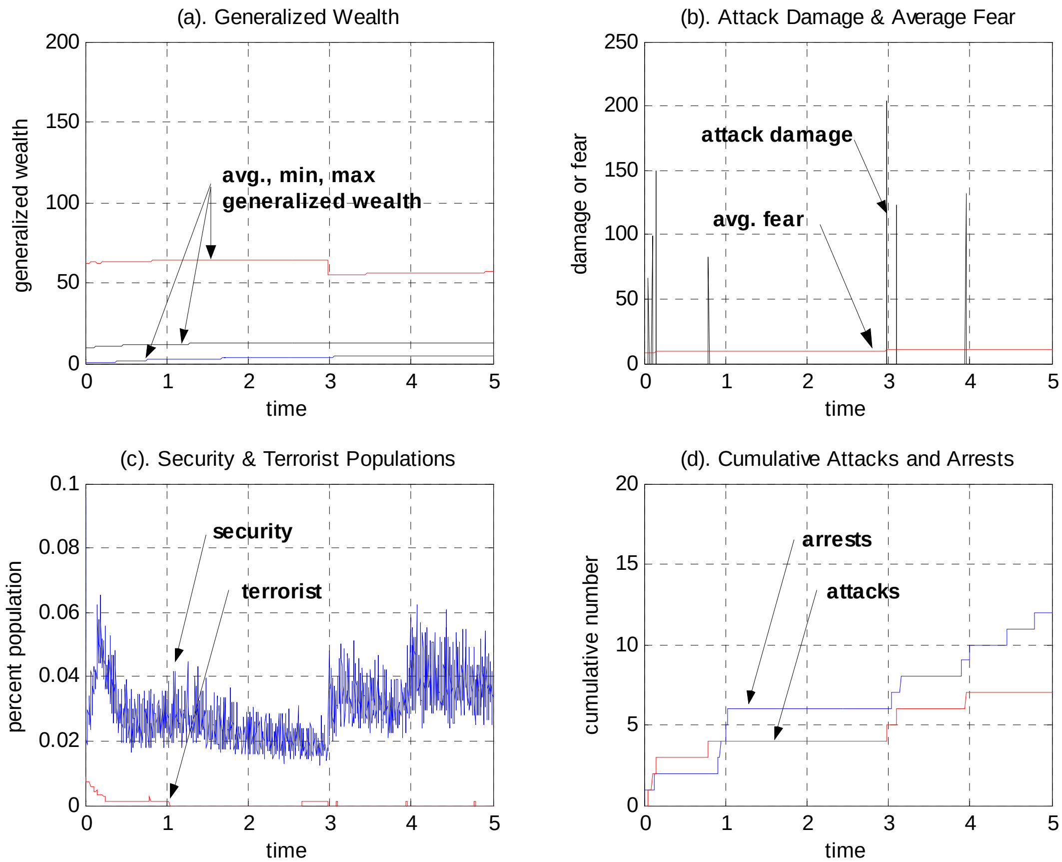

FIGURE 8 Terrorism and security activity plots for simulation 2

Further examination of Figure 8, which shows behavior typical of all four simulations, allows for some additional observations. First, all seven terrorist attacks are shown separately in Figure 8b; the attack damage represents the total agent generalized wealth destroyed by terrorists during that time-step. By comparing Figure 8b to 8c, the responsiveness of the security demand algorithm can be seen clearly. Early in the first year of the simulator, three terrorist attacks occur, and the security level spikes upward. Just prior to the end of year 1, a small terrorist attack occurs, and this attack causes a slight increase in the security level. Right around year 3, two terrorist attacks occur in quick succession; one very large one is immediately followed by another smaller one. The security level exhibits a significant increase before gradually declining. However, before the gradual decline in security can reach its pre-attack level, another terrorist attack occurs at the end of year 4. This attack causes another significant increase in security, higher than the previous jump, even though the attack was not nearly as large. The generalized wealth time history plot (average, minimum, and maximum generalized wealth for the nonsecurity population; see Figure 8a) provides a scale for the magnitude of each of the terrorist attacks relative to the agent population.

In all four simulations, there seemed to be an inordinate number of terrorist attacks occurring early in the first year. For example, for simulation run 2, Figure 8b shows that three terrorist attacks occurred in the first quarter of year one. The other simulation runs show similar scenarios. This phenomenon is primarily caused by the brief change in the behavior of civilian agents as they seek resources, which causes changes to the way the civilian agents are congregating. The terrorist agents see these increases in agent wealth concentrations and decide to conduct attacks. Once the population has settled down, terrorist behavior also settles down. This observation merely supports the previous assessment: the simulator requires a period of time to stabilize.

Effects of Changes in the Fear Memory

The sensitivity of changes in fear memory to terrorist activity and security levels was also examined. Although not directly related to simulating initial conditions, fear memory appears to be an important component in the behavior of the system. Three levels of fear memory were considered: 3-, 5-, and 8-year durations (the previous fear memory of 5 years is considered the baseline). A single incubation was run to 1,200 years for each of the fear memory cases. On the basis of the resulting post-incubation conditions, a single 5-year simulation was run for each case. All other parameters were set as previously discussed. The following discussion focuses on results from the 1,200 years of incubations. The results from the 5 years of post-incubation simulation exhibited similar trends.

First, consider the levels of terrorist and security activities. In Figures 9−11, side (a) shows the damage and average fear time history plots for the 1,200-year incubation phase for each of the three cases, and side (b) shows the cumulative number of terrorist attacks and terrorist agent arrests. A comparison of these figures indicates that the changes in fear memory have varied effects. For example, when the agents have a 5-year fear memory, 20 terrorist attacks and 11 terrorist arrests occur. When the fear memory is decreased and set at 3 years, 32 terrorist attacks and 35 terrorist arrests occur. Decreased fear memory appears to lead to an increase in terrorist and security activities. Likewise, when the fear memory is increased and set at 8 years, 82 terrorist attacks and 104 terrorist arrests occur. In other words, increased fear memory also leads to an increase in terrorist and security activities. Similar differences were observed between the corresponding 5-year simulator runs.

The relationship between fear memory and terrorist and security activity levels is the result of complex interactions within the model. Consequently, the reasons for the observed differences in the terrorist and security activity levels for the 3-, 5-, and 8-year fear memory cases can only be gleaned from a more thorough analysis of the results. For this reason, the overall historical nodal wealth should be considered. The overall historical nodal wealth is simply the average agent generalized wealth that has been present on each node over the entire incubation period of 1,200 years. Figures 12, 13, and 14 show the overall historical nodal wealth for the 3-, 5-, and 8-year fear memory cases, respectively.

Close examination of these figures will show that for the 3-year and 8-year fear memory cases, the concentration of historical nodal wealth is greater than for the 5-year fear memory case. In fact, in the 8-year fear memory case, where the terrorist and security activity levels were the highest, the concentration of historical nodal wealth was the greatest. In all three cases, the maximum, minimum, and average generalized wealth time histories were very similar; in other words, the agent populations were of similar overall wealth. These changes in the concentration of historical nodal wealth can be directly related to the terrorist agents' attack-triggering mechanism, since the value of a site includes the agent generalized wealth and historical nodal

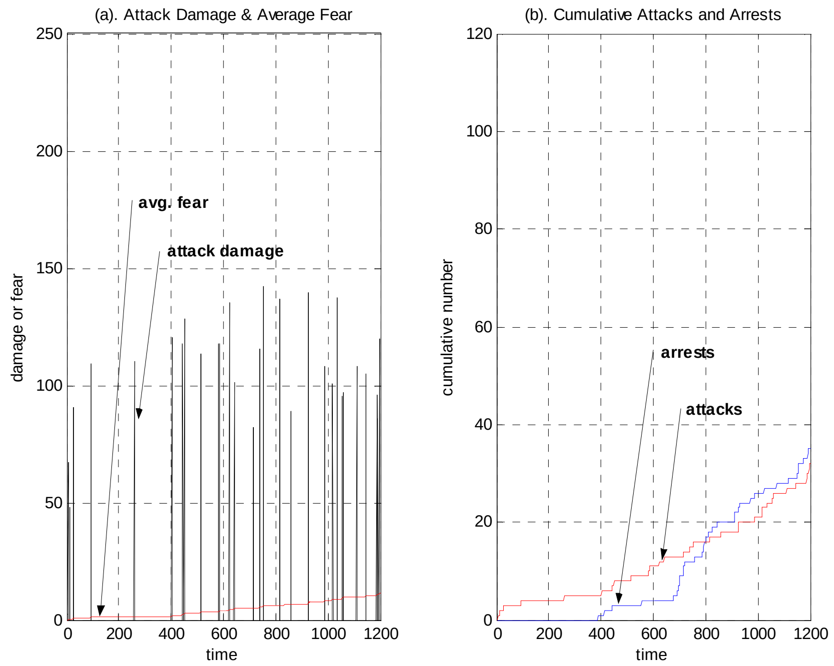

FIGURE 9 Terrorism and security statistics, 3-year fear memory

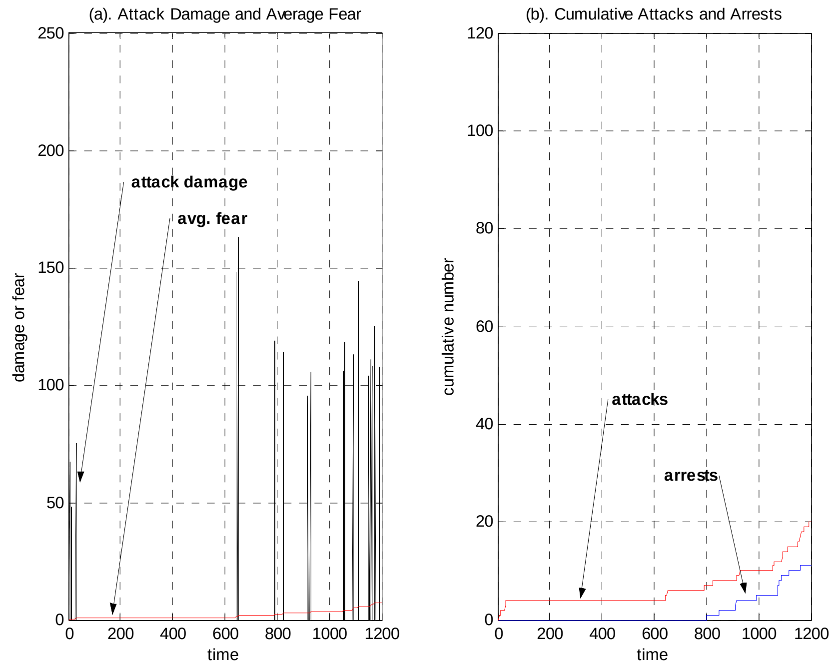

FIGURE 10 Terrorism and security statistics, 5-year fear memory

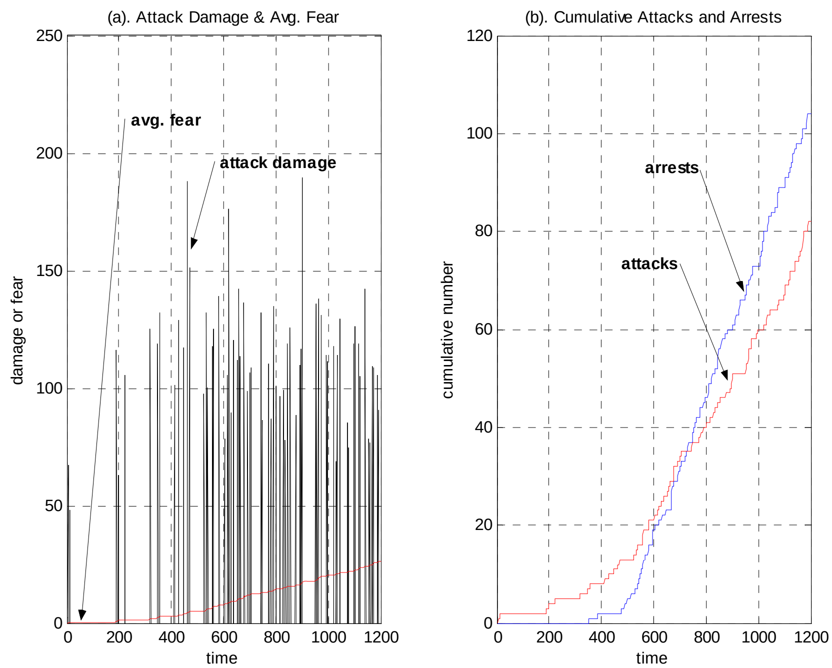

FIGURE 11 Terrorism and security statistics, 8-year fear memory

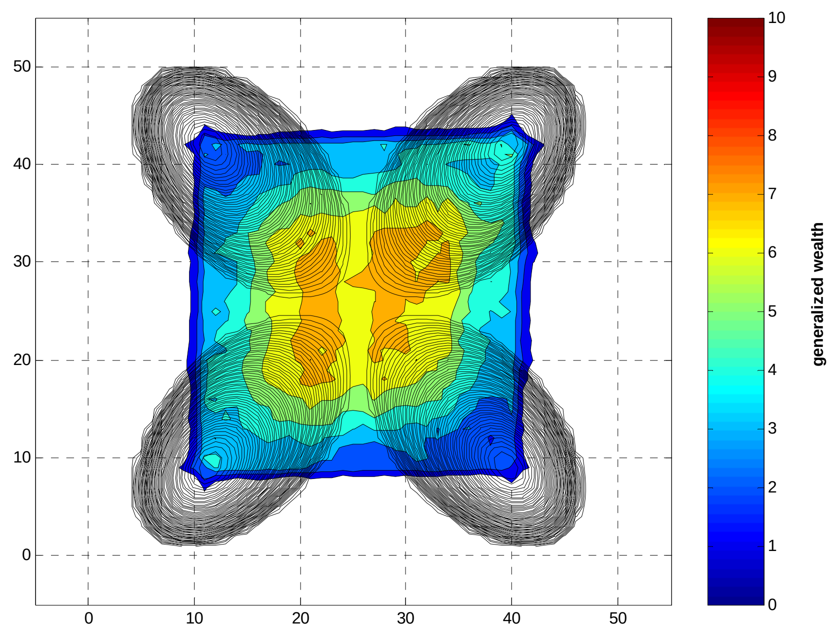

FIGURE 12 Overall historical nodal wealth, 3-year fear memory

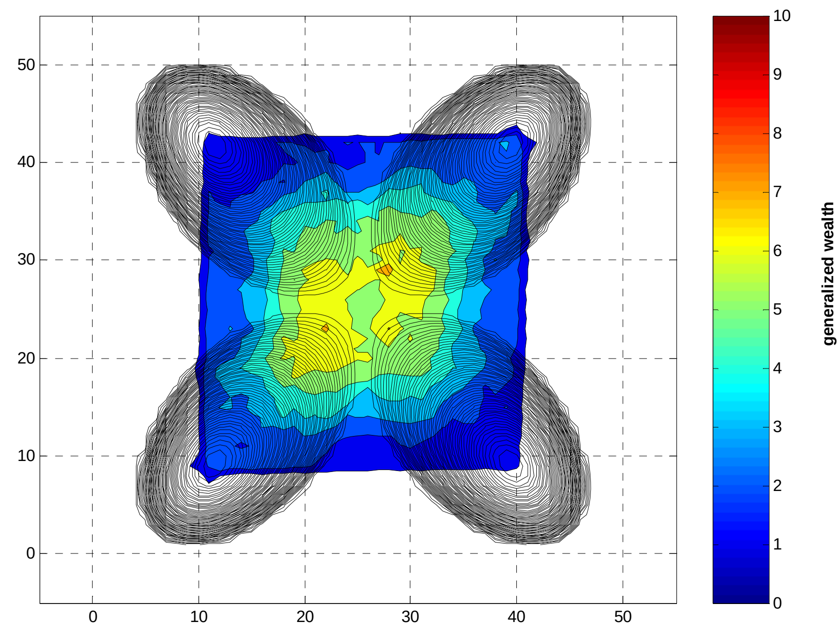

FIGURE 13 Overall historical nodal wealth, 5-year fear memory

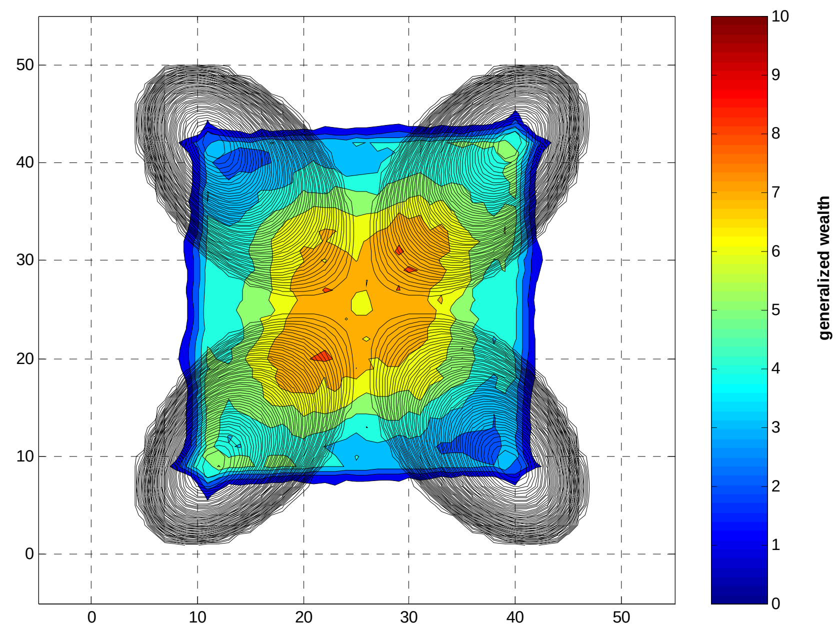

FIGURE 14 Overall historical nodal wealth, 8-year fear memory

wealth on that site's von Neumann neighborhood. The more concentrated the wealth, the more likely a terrorist attack will occur. Furthermore, if the wealth is more concentrated, the population is likely to be more concentrated too. (Overall population density contour plots verify this assumption.) The terrorist agents are more likely to be interspersed among this group and will therefore have the chance for more cultural interactions, helping to generate more agents with radical cultural identities. And finally, if more agents are in this area and it is an area where more terrorist attacks are occurring, then this will increase the security demand even further, resulting in more security agents. This is an important observation, since in the cases where the historical nodal wealth concentrations were relatively high, the security tended be more successful in arresting terrorists than terrorists were successful in carrying out attacks.

But none of this explains why different fear memory caused this increased concentration of historical nodal wealth. Recall that when civilian agents are searching for their critical resource, they are attempting to balance the reward a node has to offer (the amount of the critical resource present) with the fear to which they are subjected. Fear memory is an integral part of this balancing. The agents are looking to balance resources with fear so that they can perhaps avoid becoming a victim of a terrorist attack. However, there is apparently an optimal memory. Too great a fear memory leads agents to behave in a way that makes them become a prime target for a terrorist attack. Too little of a fear memory has agents moving around, oblivious to the risk; therefore, agents are again behaving in a way that makes them become a prime target for a terrorist attack. In other words, fear memory is a tool that can either help or hurt, depending on its scope. Further work is required to confirm that this effect of fear memory is a general trend and not just an artifact of this particular case.

CONCLUSIONS

A resource-based agent model has been developed to model terrorist activity. Endogenous terrorist agents are formed from within the civilian agent population by using a tag-mediated cultural identity. The terrorist agents conduct surveillance and commit terrorist attacks. When a terrorist attack occurs, fear is generated in the area subjected to the attack. Over time, this fear spreads out to the surrounding area. When searching for resources, civilian agents attempt to balance the rewards of visiting certain nodes with their fear of those nodes and with their innate nervousness. Civilian agents demand security on the basis of their wealth, fear, and nervousness.

When the model is run by using a time-step analogous to 1 year, the model is said to be acting as an incubator. After the incubator is run for a period of time, the agent population evolves from that set by the user to a population in tune with its environment. When the incubation phase ends, the post-incubation conditions become the initial conditions for a relatively short time-step simulation: the simulator.

A number of incubator and simulator runs were conducted. First the incubator was run and terminated at various points, thereby generating post-incubation conditions at various stages of the population's development. Second, multiple post-incubation conditions were generated by using identical input parameters. On the basis of these results, the following conclusions could be drawn. (1) In the transition between incubation and simulation, the characteristics defining the agent population remain stable. (2) The level of terrorist activity and, hence, the level of security in the environment remain consistent between the incubator and simulator. (3) The incubator and simulator are robust, requiring no special criteria to be satisfied in order to generate initial conditions. (4) The behavior of civilian agents when searching for resources is altered when they enter the simulator but stabilizes after a period of time (the length of which depends on the level of evolution attained in the incubator).

A combined incubator/simulator simulation was performed for a 3-, 5-, and 8-year fear memory. The 5-year fear memory case was considered as the baseline. The following qualitative conclusions resulted. (1) An increase in the duration of the fear memory resulted in significant increases in the level of terrorist activity and consequently the security level. (2) A decrease in the duration of fear memory also resulted in an increase in terrorism and a larger security presence. (3) The increased terrorism seems to result from changes in civilian agent behavior — changes that lead to increased concentrations of wealthy agents. (4) The changes in the behavior of civilian agents may be allowing increased numbers of radical agents to form (increased population densities lead to more cultural interactions) and may be making it easier for terrorists to conduct attacks. Thus, the duration of fear memory (i.e., the amount of time that an agent remembers fear that it has seen) can have significant effects on the level of terrorist activity. The effect is nonlinear; too great or too little fear memory works against the civilian agents, promoting increases in terrorism. This effect must be examined more carefully.

ACKNOWLEDGMENTS

We would like to thank both the National Science Foundation Graduate Research Fellowship Program and the Michigan Technological University Department of Civil and Environmental Engineering. Their support has made this project possible. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the writers and do not necessarily reflect the view of the National Science Foundation.

REFERENCES

Axelrod, R., 1997, *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*, Princeton, NJ: Princeton University Press.

Bulleit, W.M., and M.W. Drewek, 2005, “An Agent-based Model of Terrorist Activity,” in *Proceedings of the North American Association for Computational Social and Organizational Science (NAACSOS) Conference,* on CD, Notre Dame, IN, June 26–28.

Cederman, L.-E., 2003, “Modeling the Size of Wars: From Billiard Balls to Sandpiles,” *American Political Science Review* 97(1):135–150.

Dorfman, R., and N.S. Dorfman, 1993, *Economics of the Environment: Selected Readings,* New York, NY: W.W. Norton & Co.

Epstein, J.M., 2002, “Modeling Civil Violence: An Agent-based Computational Approach,” in *Proceedings of the National Academy of Sciences*, Vol. 99, Suppl. 3; available at http://www.pnas.org/cgi/doi/10.1073/pnas.092080199.

Epstein, J.M., and R. Axtell, 1996, *Growing Artificial Societies: Social Science from the Bottom up*, Washington, DC: The Brookings Institution.

MathWorks, 2002, *MatLab Student Version with Simulink*, Version 6.5, Release 13, Natick, MA.